



FAX TRANSMISSION

UNIVERSITY OF CAMBRIDGE

Department of History and
Philosophy of Science
Free School Lane
Cambridge CB2 3RH
(international)

Fax No. 01223 334554 (national)
+44 1223 3345540 (international)
Tel. 01223 334540 (national)
+ 44 1223 334540

To: *Prof. Giancarlo Ghisardi*
Dept of Theoretical Physics
University of Trieste

Fax No: *+39 40 224601*

From: *Prof. M. L. S. Redhead*

Number of pages including this one *20* Date *23/01/96*
Message:

If you do not receive legible copies, please notify us at once

UNIVERSITY OF CAMBRIDGE
DEPARTMENT OF HISTORY AND PHILOSOPHY OF SCIENCE

FREE SCHOOL LANE
CAMBRIDGE CB2 3RH
TEL : 01223 334540(334556)
FAX: 01223 334554

Professor Michael Redhead FBA
Chairman of Department

21.1.96

Dear Gian Carlo,

I have now had time to return to your letter of 6th December and would like to offer the following comments.

There are two crucially distinct readings of OM-LOC:

(1) There is the reading which prohibits measurement procedures on the left from affecting measurement outcomes on the right. This is a case by case version of the probabilistic parameter independence, P.I.

This is the reading which is actually used in the derivation of
 $B.M. \wedge \text{Compl.} \wedge (OM-LOC) \supset \neg (ER-LOC) *$
P.T.O.

and hence the denial of OM-Loe in this sense can indeed block the derivation of $\neg(ER-Loe)$ as we all agree.

(2) But there is a logically stronger sense that prohibits not only setting-to-result effects at space-like separation, but also result-to-result effects (i.e. a case by case version of O.I.)

Now it is crucial for the logical structure of your argument that you use the logically weakest version of OM-Loe that is necessary for the result x to hold.

If you use a stronger version than is necessary, you cannot infer that a failure of this stronger version will block the derivation of $\neg(ER-Loe)$.

That is just a matter of straight logic.

Now in your discussion of the relation between OM-Loe and B-Loe, you are all

the time using ^{II} the stronger version
of OM-loc, since in your examples
you refer to the question of whether
the outcomes of the two measurements
influence one another.

Referring to my previous letter
dated 5th November, I was always
using the weak version of OM-loc,
and for that version, I maintain
that my discussion of the
problematics ~~of~~ of deriving

$$\text{OM-loc} \supset \text{B-loc}$$

was correct.

I would like now to deal with your
Example 3, in which you claim that
 $\text{OM-loc} \supset \text{P.I.}$, requires a No Conspiracy
Assumption.

If Alice selects ^{by some rule} ~~when~~ to toss, then she
is inducing a place-selection on the
random sequence generated by Bob. If
this altered the limiting frequency of Bob's
results, this would contradict the randomness
P.I.V

of Bob's sequence. But if Bob's sequence really is random, this could not happen!

If you agree with the points in this letter, then I suggest the best way forward would be as follows: I enclose a copy of a revised version of the paper with La Rivière, which I think summarizes my considered point of view on this vital question of a correct relativistic formulation of the EPR argument. I have been asked by Bob Cohen whether I would consider contributing this paper to the Shimony Festchrift. He is editing. I would be happy to do this, and then you could raise additional points and comments in a further submission to Studies in History and Philosophy of Modern Physics.

I would play no part in the editorial discussion of such a submission,

III

but would leave this for Jeremy to sort
out with you.

I have very much enjoyed, and
benefited from, our discussions of this
matter.

With very best wishes

Michael

P.S. I should add in relation to
your discussion of why the violation
of OM-Loc is less serious than
the violation of ER-loc from a
relativistic point of view, that
you again use the strong version
of OM-Loc, when you talk, for
example, of a "violation of OM-Loc -- due
to a violation of O.I."

The Relativistic EPR Argument:
A Critique of Ghirardi and Grassi

Michael Redhead and Patrick La Rivière

Department of History and Philosophy of Science
Cambridge University

An obvious drawback of realist interpretations of quantum mechanics such as Bohm's is their nonlocal character, which seems to conflict with the constraints of relativity theory. These interpretations seem no worse in this regard, however, than the standard 'antirealist' interpretations that reject hidden variables and assume the completeness of the quantum formalism, because the original Einstein-Podolsky-Rosen argument entails nonlocality in these cases as well. But the standard EPR argument is not formulated in a relativistic context—obviously a crucial shortcoming if we want to compare these realist and antirealist interpretations on the basis of their agreement with relativity theory. We thus propose in this paper to investigate the possibility of a relativistic formulation of the EPR argument. In particular, we will examine the recent attempt by G. Ghirardi and R. Grassi to formulate such a relativistic EPR argument, from which they conclude that "the coexistence of quantum mechanics with relativity is even more peaceful than is commonly thought." The present paper will criticize the argument employed by Ghirardi and Grassi as being derived from an ambiguously stated locality principle. However, it is argued further that the reasoning used by Ghirardi and Grassi employs a problematic interpretation of counterfactuals, which under an assumption of indeterminism potentially undermines the whole line of argument while, ironically, salvaging their original conclusion. This conclusion, naturally, offers a point in favour of the antirealist interpretations over the realist views.

1. Introduction

At first glance, the realist interpretations of quantum mechanics such as Bohm's offer many advantages over standard interpretations of the theory. In particular, they give a clear, intuitive picture of many potentially paradoxical physical situations, such as the two-slit experiment and the phenomenon of barrier penetration. At the same time, their chief drawback—a form of nonlocality that seems to conflict with the constraints of relativity theory—is apparently shared by the standard, 'antirealist' interpretations that reject hidden variables and assume completeness, as was demonstrated by the original Einstein-Podolsky-Rosen argument. However, while the Bell argument that establishes nonlocality for realistic interpretations such as Bohm's has been formulated in a relativistic context (Landau 1987; Summers and Werner 1985), there is no well-established relativistic formulation of the EPR argument. In the absence of such a formulation, it seems hasty to conclude that the tension between the standard interpretations and relativity theory is just as great as that between Bohmian interpretations and relativity. Clearly, if a relativistic formulation of EPR could be given that did *not* entail nonlocality, antirealist interpretations would have an advantage over the Bohmian interpretation.

The present paper will investigate the possibility of a relativistic formulation of the EPR argument. In Section 2, we will review the standard non-relativistic version of the EPR argument and consider the problematics of translating it into a relativistic context. We will pay particular attention to the need for a reformulation of the so-called reality criterion. In section 3, we introduce one such reformulated reality criterion, due to Ghirardi and Grassi (1994), and we show how it is applied to the non-relativistic EPR argument. In section 4, we discuss the application of the new reality criterion in a relativistic context and point out a flaw in Ghirardi and Grassi's argument that undermines their conclusion of peaceful co-existence. Finally, in section 5,

we engage issues related to the evaluation of counterfactuals that reveal a hidden assumption of determinism in Ghirardi and Grassi's proof, while offering a way of salvaging their conclusion. In section 6, we review and summarize our own conclusions.

2. Non-Relativistic EPR

A relativistic version of the EPR argument must differ from the non-relativistic version in two principal ways. First, the particle states must be described by a relativistic wavefunction. The details don't concern us here; we need only require that the wavefunction preserve the maximal, mirror-image correlations of the non-relativistic singlet state. And indeed, one of us¹ has recently demonstrated the existence of maximal correlations in the vacuum state of relativistic algebraic quantum field theory. Second, the argument must not depend on the existence of absolute time ordering between the measurement events on the left and right wings of the system, for in the relativistic argument these may be space-like separated. As it turns out, the non-relativistic version of the argument *does* invoke absolute time ordering. To see how to get around this problem, we must briefly review the standard formulation of the incompleteness argument.

For EPR, a necessary condition for the completeness of a theory is that every element of physical reality have a counterpart in the theory. To demonstrate that quantum mechanics is incomplete, EPR need simply point to an element of physical reality that does not have a counterpart in the theory. In this vein, they consider measurements on a pair of scattered particles with correlated position and momentum, but Bohm's formulation of the argument (1951), in terms of a pair of oppositely moving, singlet-state, spin-1/2 decay products of a spin-0 particle, is conceptually simpler. In this case, the formalism of quantum mechanics demands a strict correlation between the spin components of the two spatially separated particles, such

that a measurement of, say, the z-component of spin of one particle allows one to predict with certainty the outcome of the same measurement on the distant particle. This ability to predict with certainty, or at least probability one, the outcome of a measurement is precisely the EPR criterion for the existence of an element of reality at the as-yet-unmeasured particle. By invoking one final assumption, a locality assumption stating that elements of reality pertaining to one system cannot be affected by measurements performed 'at a distance' on another system, EPR can establish that the element of reality at the unmeasured particle must have existed even *before* the measurement was performed at the distant particle. But the quantum-mechanical formalism describes the particles at this point with the singlet state, and thus has no counterpart for the element of reality at the unmeasured particle. It follows that the quantum-mechanical description was incomplete.² Schematically,

$$\text{QM Formalism} \wedge \text{Locality} \rightarrow \sim \text{Completeness}$$

Alternatively, if one assumes completeness, the argument may be rearranged as a proof of nonlocality:

$$\text{QM Formalism} \wedge \text{Completeness} \rightarrow \sim \text{Locality}$$

The problematic assumption of absolute time ordering entered the argument in the reality criterion, which turns on the possibility of predicting with certainty the outcome of a measurement along one wing *subsequent* to having obtained the result of a measurement along the other. Of course, for space-like separated events, notions like precedence and subsequence are reference-frame dependent, not absolute. So to translate the EPR argument to a relativistic context requires a new criterion for the attribution of elements of

reality that is not contingent on the time ordering of the measurement events. In a recent paper, Ghirardi and Grassi (1994) have undertaken to formulate just such a criterion, and thus to salvage the EPR argument in a relativistic framework. For the sake of clarity, we shall first describe how this criterion applies to the non-relativistic version of the argument.

3. The Counterfactual Criterion and Non-Relativistic EPR

Ghirardi and Grassi's criterion rests on the truth of certain classes of counterfactual statements³—statements of the form 'if ϕ were true, then ψ would be true', where the antecedent ϕ is in general known to be false. In particular, they wish to 'link...the attribution at time t of the property corresponding to [observable A having value a] to the truth of the counterfactual assertion: if a measurement of $[A]$ were performed at time t , then the outcome would be $[a]$.' In order to evaluate the truth of such statements, they call on the work of David Lewis (1973):

Let us denote the counterfactual 'if ϕ were true then ψ would be true' as ' $\phi \Box \Rightarrow \psi$ '. Then Lewis proposes the following truth condition: ' $\phi \Box \Rightarrow \psi$ ' is true at world w if either (i) there are no possible worlds at which ϕ is true or (ii) some world where both ϕ and ψ are true is more similar ('closer') to w than any world in which ϕ is true and ψ is false. Obviously one has to specify the possible worlds one is taking into account; this is done by assigning to each world w a set of worlds S_w called the sphere of accessibility around w .

With this criterion in hand, Ghirardi and Grassi can now run the *non-relativistic* EPR argument essentially as before. They assume a measurement of property A is performed on the right-hand system at time t_R , yielding a specific result a . To ascertain whether an element of reality corresponding to property $A = a'$ exists at the left-hand system, they must assess the truth of the counterfactual assertion: 'if I were to perform a measurement of property

A at the left-hand system at time t_L , I would obtain the result a' .' In the non-relativistic case, the truth of this counterfactual assertion follows naturally from the presence of absolute time ordering. For if $t_R < t_L$, then the outcome of the right-hand measurement can be assumed to be the same in all of the 'accessible' (most similar) worlds used to evaluate the counterfactual, because it is strictly in the past of the counterfactual's antecedent. The strict correlation laws of quantum mechanics, also assumed to hold in all accessible worlds, then demand that the result of a measurement on the left wing also be fixed in all possible worlds (specifically, the laws require that $a' = -a$). Thus the counterfactual is true, and an element of reality can be said to exist at the left-hand system. From here, the argument unrolls in the usual way, and by supplementing this reality criterion with a locality assumption (they call it G-Loc, after Galileo), Ghirardi and Grassi can deduce that quantum mechanics is incomplete. Once again, we can represent their argument schematically by

$$\text{QM Formalism} \wedge \text{G-Loc} \rightarrow \sim \text{Completeness}$$

or

$$\text{QM Formalism} \wedge \text{Completeness} \rightarrow \sim \text{G-Loc}$$

While these conclusions seem sound, the locality principle, G-Loc, bears further investigation. It reads: "A system cannot be affected by actions on a system from which it is isolated. In particular, elements of physical reality of a system cannot be influenced by actions on systems from which it is isolated." An examination of the structure of Ghirardi and Grassi's argument reveals that they make use not of the general principle stated but of a special case of this general principle, viz. that elements of reality cannot be brought into existence 'at a distance.' It is this special case of G-Loc, call it ER-Loc (for

elements of reality) that enters toward the end of the argument to establish that the measurement at the right wing could not have created an element of reality at the left wing and thus that it must have existed prior to the measurement at the right wing, when the quantum formalism said the particles were in the singlet state. Thus they conclude that quantum mechanics is incomplete. All is well so far, but when one turns the argument around, assuming completeness and dispensing with locality, one must ask, can one be more precise as to which locality principle should be given up: the principle they label G-Loc, or the special case ER-Loc? Indeed it is the latter, for only it entered into the argument. As it turns out the distinction between G-Loc and ER-Loc does not affect their conclusions in the non-relativistic case, because the conclusion they choose to highlight—the creation of elements of reality at a distance—is precisely one that does follow from dispensing only with ER-Loc.

4. Relativistic EPR⁴

In the relativistic case, however, more care must be taken with the statement of the locality principle, this time called L-Loc (after Lorentz), because a locality principle must enter at the very beginning of the argument as well as in the usual way at the end. The argument begins in the same way as in the non-relativistic case, with the occurrence of a measurement on the right-hand side, but now the absence of absolute time ordering means the result of this measurement can no longer tacitly be assumed to be the same in all the accessible worlds used to evaluate the element-of-reality counterfactual at the left-hand side. Locality must be invoked to establish the independence of the outcome of the right-hand measurement from the occurrence of the measurement at the left. This done, Ghirardi and Grassi then demonstrate the existence of an element of reality at the left-hand side following the same reasoning as above. From here, the argument unrolls once again in the usual

way and locality makes a second appearance in its familiar place at the end of the argument. In this way, Ghirardi and Grassi can again prove that standard quantum mechanics plus 'locality' implies incompleteness.

There are two quite distinct cases of L-Loc that are actually being employed, one used in getting the argument started and the other serving as the conclusion. Ghirardi and Grassi define L-Loc as the following: "An event cannot be influenced by events in space-like separated regions. In particular, the outcome obtained in a measurement cannot be influenced by measurements performed in space-like separated regions; and analogously, possessed elements of physical reality referring to a system cannot be changed by actions taking place in space-like separated regions." As in the non-relativistic case, it is not the general principle but rather the two special cases, call them OM-Loc (for outcome of measurement) and ER-Loc (again for elements of reality), that are doing the logical work in their argument. They invoke OM-Loc at the beginning of the argument while assessing the counterfactual reality criterion, as discussed above, and they invoke ER-Loc at the end of the argument, as they did in the non-relativistic case. So if we write $L\text{-Loc} = OM\text{-Loc} \wedge ER\text{-Loc}$, then, schematically, their argument looks like this:

$$\text{Quantum Formalism} \wedge OM\text{-Loc} \wedge ER\text{-Loc} \rightarrow \sim \text{Completeness}$$

or

$$(*) \quad \text{Quantum Formalism} \wedge \text{Completeness} \rightarrow \sim OM\text{-Loc} \vee \sim ER\text{-Loc}$$

Ghirardi and Grassi now argue effectively as follows. Assuming OM-Loc we can again demonstrate from Completeness a violation of ER-Loc, i.e. Einstein's *spooky* action-at-a-distance *creating* elements of reality at a distance. But if we don't assume OM-Loc, then we cannot deduce a violation of ER-Loc. All this is quite correct, but the price we have to pay for *not* being able to

demonstrate a violation of ER-Loc is precisely that we have to accept a violation of OM-Loc!

In other words the relativistic formulation of the EPR argument does not help with the thesis of peaceful coexistence between quantum mechanics and special relativity, unless one argues that violating ER-Loc is more serious than violating OM-Loc from a relativistic point of view.

If we rewrite (*) as

$$\text{Quantum Formalism} \wedge \text{Completeness} \wedge \text{L-Loc} \rightarrow \sim \text{L-Loc}$$

it would be easy to misconstrue Ghirardi and Grassi as illegitimately claiming that (*) cannot be used to demonstrate $\sim \text{L-Loc}$ on the grounds that L-Loc is presupposed in the argument. This is *not* what they say, but to justify peaceful coexistence, we need to identify an additional assumption omitted from (*), which, if challenged, could undermine the inference. This we shall proceed to investigate in the next section.

5. Counterfactuals and Indeterminism

Recall that to run the argument in either the non-relativistic or relativistic case, Ghirardi and Grassi must establish that the outcome of, say, the right-hand measurement is the same in all accessible worlds. With this established, the correlation laws of quantum mechanics imply that the outcome of the left-hand measurement is the same in all accessible worlds, and hence establish the truth of the counterfactual assertion about the left-hand measurement result that permits the attribution of an element of reality to the left-hand system. In the non-relativistic case, the constancy of the right-hand result follows trivially from absolute time ordering as discussed above; in the relativistic case, it's not so simple. A premise akin to one that Redhead,

following Stapp, labels the Principle of Local Counterfactual Definiteness (PLCD) is needed to do this sort of work (Redhead 1987, 92).

In the present case, PLCD may be taken to assert that the result of an experiment which *could* be performed on a microscopic system has a definite value that does not depend on the occurrence of a measurement at a distant apparatus. Ghirardi and Grassi implicitly assume that PLCD is licensed by their locality principle, for they invoke only OM-Loc to establish the constancy of the right-hand outcome in all accessible worlds. But Redhead argues that PLCD does *not* follow directly from any typical locality principle, certainly not from one like OM-Loc, which asserts that the outcome obtained in a measurement cannot be influenced by measurements in space-like separated regions. The reason is quite simple: while invoking locality may prevent measurements on the left-hand system from influencing the result at the right and breaking the constancy of the accessible worlds as far as the right-hand result is concerned, it does not prevent *indeterminism* from wreaking that sort of havoc. Intuitively, we can imagine that we run the world over again, this time performing the measurement on the left-hand system. If we consider this left-hand measurement schematically as a point event with a backward light cone identical to that in the actual world, we are concerned with what will happen in the complement of the forward and backward light cones. Under indeterminism, we claim, the events in this complement simply *cannot* be assumed to remain the same. This claim is not uncontroversial, however, for Lewis himself has argued that the events in the complement can be assumed to be fixed; thus for Lewis, OM-Loc *does* licence PLCD.⁵ His argument turns on a dual reading of the 'might' counterfactual implicit in our description of 're-running' the world: if I were to run the world over again and perform the left-hand measurement, the right-hand outcome *might* be different than it was in the actual world. This 'might,' he argues, could be read either as 'would be possible' or as 'not would not,' but that the

first reading does not contradict the negation of the second reading. We are not enamoured of this slippery semantic solution to the problem and thus maintain that Ghirardi and Grassi need both OM-Loc and an assumption of determinism to get their argument off the ground. Schematically, (*) is replaced by

$$(**) \quad \text{Quantum Formalism} \wedge \text{Completeness} \wedge \text{Determinism} \rightarrow \\ \sim \text{OM-Loc} \vee \sim \text{ER-Loc}$$

It seems, then, that Ghirardi and Grassi's reformulation of the EPR argument in a relativistic context is less general than they would have us believe, for its scope is limited to deterministic systems.⁶

6. Conclusion

We have examined Ghirardi and Grassi's attempt to reformulate the EPR argument in a relativistic context and found it to be flawed by an ambiguously stated locality principle and a hidden assumption of determinism. By making explicit the logical structure of their argument, we have undermined the conclusion that in the relativistic case the existence of action-at-a-distance is not a valid deduction from the EPR argument.

This conclusion can only be rescued if an additional hidden assumption of determinism is exposed. Assuming indeterminism then, we can avoid the EPR inference to action-at-a-distance, and the concomitant challenge to peaceful co-existence between quantum mechanics and special relativity. Thus we end up agreeing with Ghirardi and Grassi, but for different reasons than the ones they present in their paper.

¹ See Redhead (1995). The argument actually applies more generally to any state of bounded energy; see also Redhead (1994).

² We follow here a streamlined version of the EPR argument, as introduced by Redhead (1987) and Hellman (1987). Historically this version of the argument seems to have been known to Einstein. See, for example, Fine (1986). This is also the version used by Ghirardi and Grassi in the paper under discussion.

³ An analysis of the EPR argument using counterfactuals, though not specifically in the context of a relativistic reformulation, has been undertaken by Wessels (1981). She seeks to uncover the full logical structure of the EPR argument by formulating the original paper's somewhat ambiguous reality criterion in precise modal terms. Among four possible modal readings of the EPR reality criterion, she lists a counterfactual reading similar to Ghirardi and Grassi's, which is the one we also adopt in this paper.

⁴ A discussion of the EPR set-up in a relativistic context has also been provided by G. Smith and R. Weingard (1987). They argue that any relativistic formulation of EPR should employ a relativistic correlated state. They derive such a state and demonstrate the relativistic invariance of the correlations. However they fail to pursue the analysis beyond the existence of the correlations, i.e. to develop the full EPR argument. Three-particle versions of the EPR argument in a relativistic context have also been considered recently in the literature. See, for example, Clifton, Pagonis, and Pitowsky (1992). For a critique of their conclusions, see Pagonis, Redhead, and La Rivière (1995).

⁵ D. Lewis, private communication. See also Lewis (1986).

⁶ This whole topic of evaluating counterfactuals under an assumption of indeterminism has been the subject of a long-running debate between H. Stapp and M. Redhead. See in particular Clifton, Butterfield, and Redhead (1990). The most up-to-date list of references on the topic can be found in Clifton and Dickson (1994) and Stapp (1994).

References

- Bohm, D. (1951), *Quantum Theory*. Englewood Cliffs, Prentice-Hall.
- Clifton, R., Butterfield, J., and Redhead, M. L. G. (1990), "Nonlocal Influences and Possible Worlds—A Stapp in the Wrong Direction," *British Journal for the Philosophy of Science* **41**, 5–58.
- Clifton, R. and Dickson, M. (1994), "Stapp's Algebraic Argument for Nonlocality," *Physical Review A* **49**, 4251–4256.
- Clifton, R., Pagonis, C., and Pitowski, I. (1992), "Relativity, Quantum Mechanics and EPR," in D. Hull, M. Forbes and K. Okruhlik (eds.), *PSA 1992*, vol 1. East Lansing, Philosophy of Science Association, pp. 114–128.
- Fine, A. (1986), *The Shaky Game*. Chicago, University of Chicago Press.
- Ghirardi, G. and Grassi, R. (1994), "Outcome Predictions and Property Attribution: the EPR Argument Reconsidered," *Studies in the History and Philosophy of Science* **25**, 397–423.
- Hellman, G. (1987), "EPR, Bell and Collapse: A Route Around 'Stochastic' Hidden Variables," *Philosophy of Science* **54**, 558–576.
- Pagonis, C., Redhead, M. L. G., and La Rivière, P. (1995), "EPR, Relativity, and the GHZ Experiment," forthcoming in R. Clifton (ed.), *Perspectives on Quantum Reality*. Dordrecht, Kluwer.
- Landau, L. J. (1987), "On the Non-Classical Structure of the Vacuum," *Physics Letters A* **123**, 115–118.
- Lewis, D. (1973), *Counterfactuals*. Oxford, Basil Blackwell Ltd.
- Lewis, D. (1986), *Philosophical Papers*, vol. 2. Oxford, Oxford University Press.
- Redhead, M. L. G. (1987), *Incompleteness, Nonlocality, and Realism*. Oxford, Oxford University Press.
- Redhead, M. L. G. (1994), "The Vacuum in Relativistic Quantum Field Theory," forthcoming in D. Hull, M. Forbes and K. Okruhlik (eds.), *PSA 1994*, vol. 2. East Lansing, Philosophy of Science Association.
- Redhead, M. L. G. (1995), "More Ado About Nothing," *Foundations of Physics* **25**, 123–137.

- Smith, G. and Weingard, R. (1987), "A Relativistic Formulation of the Einstein-Podolsky-Rosen Paradox," *Foundations of Physics* 17, 149–172.
- Stapp, H. (1994), "Reply to 'Stapp's Algebraic Argument for Nonlocality,'" *Physical Review A* 49, 4257–4260.
- Summers, S.J. and Werner, R. (1985), "The Vacuum Violates Bell's Inequalities," *Physics Letters A* 110, 257–259.
- Wessels, L. (1981), "The EPR Argument: A Post-Mortem," *Philosophical Studies* 40, 3–30.